

Moment based gene set tests

Jessica L. Larson
Genentech Inc.

Art B. Owen
Stanford University

April 2014

Abstract

Motivation: Permutation-based gene set tests are standard approaches for testing relationships between collections of related genes and an outcome of interest in high throughput expression analyses. Using M random permutations, one can attain p -values as small as $1/(M+1)$. When many gene sets are tested, we need smaller p -values, hence larger M , to achieve significance while accounting for the number of simultaneous tests being made. As a result, the number of permutations to be done rises along with the cost per permutation. To reduce this cost, we seek parametric approximations to the permutation distributions for gene set tests.

Results: We focus on two gene set methods related to sums and sums of squared t statistics. Our approach calculates exact relevant moments of a weighted sum of (squared) test statistics under permutation. We find moment-based gene set enrichment p -values that closely approximate the permutation method p -values. The computational cost of our algorithm for linear statistics is on the order of doing $|G|$ permutations, where $|G|$ is the number of genes in set G . For the quadratic statistics, the cost is on the order of $|G|^2$ permutations which is orders of magnitude faster than naive permutation. We applied the permutation approximation method to three public Parkinson's Disease expression datasets and discovered enriched gene sets not previously discussed. In the analysis of these experiments with our method, we are able to remove the granularity effects of permutation analyses and have a substantial computational speedup with little cost to accuracy.

Availability: Methods available as a Bioconductor package, npGSEA (www.bioconductor.org).

Contact: larson.jessica@gene.com

1 Introduction

In a genome-wide expression study, researchers often compare the level of gene expression in thousands of genes between two treatments groups (e.g., disease, drug, genotype, etc.). Many individual genes may trend toward differential expression, but will often fail to achieve significance. This could happen for a set of genes in a given pathway or system (a gene set). A number of significant

and related genes taken together can provide strong evidence of an association between the corresponding gene set and treatment of interest. Gene set methods can improve power by looking for small, coordinated expression changes in a collection of related genes, rather than testing for large shifts in many individual genes.

Additionally, single gene methods often assume that all genes are independent of each other; this is not likely true in real biological systems. With known gene sets of interest, researchers can use existing biological knowledge to drive their analysis of genome-wide expression data, thereby increasing the interpretability of their results.

Mootha et al. (2003) first introduced gene set enrichment analysis (GSEA) and calculated gene set p -values based on Kolmogorov-Smirnov statistics. Since then, there have been many methodological proposals for GSEA; no single one is always the best. For example, some tests are better for a large number of weakly associated genes, while others have better power for a small number of strongly associated genes (Newton et al., 2007).

One of the most important differences among gene set methods is the definition of the null hypothesis. Tian *et al.*, 2005 and Goeman and Bühlmann, 2007 (among others) introduce two null hypotheses that differentiate the general approaches for gene set methods. The first measures whether a gene set is more strongly related with the outcome of interest than a comparably sized gene set. Methods of this type typically rely on randomizing the gene labels to test what is often called the *competitive* null hypothesis. This is problematic because genes are inherently correlated (especially those within a set) and permuting them does not give a rigorous test (Goeman and Bühlmann, 2007).

The second type of approach is used to determine whether the genes within a set associate more strongly with the outcome of interest than they would by chance, had they been independent of the outcome. Methods that test this *self-contained* null hypothesis usually judge statistical significance by randomizing the phenotype with respect to expression data and assume that gene sets are fixed. While we acknowledge that the *competitive* hypothesis is often of interest, we focus on methods that test the *self-contained* hypothesis in this paper.

A popular *self-contained* GSEA method is the JG-score (Jiang and Gentleman, 2007), which determines the level of enrichment based on averaging linear model statistics. Recently, Ackermann and Strimmer (2009) compared 261 different gene set tests, and found particularly good performance from a sum of squared single gene regression coefficients. We extend both the sum and the sum of squared linear statistics approaches with a new method in this paper.

All current GSEA methods are based on permutation approaches. The initial GSEA (Mootha et al., 2003) and JG-score (Jiang and Gentleman, 2007) methods both have closed form null distributions for their enrichment statistics, Gaussian and Kolmogorov-Smirnov, respectively; however, even the authors of these methods acknowledge that these distributions do not give the correct p -values and suggest the use of permutation. Lehmann and Romano (2005) give a concise explanation of how permutation inference works. It is common to approximate the permutation distribution by a large Monte Carlo sample (Eden

and Yates, 1933; David, 2008).

Permutation tests are simple to program and do not make parametric distributional assumptions. They also can be applied to almost any statistic we might wish to investigate. However, permutation approaches are often computationally expensive, are subject to random inference, and fail to achieve continuous p -values. Each of these drawbacks is described in more depth below.

We have developed a new gene set enrichment approach that approximates the permutation distribution of our corresponding test statistics. We find that our method of moments techniques result in almost exactly the same p -values as permutation approaches, but in much less computation time. Through our approach, we are able to obtain refined p -values and achieve stringent significance thresholds. We applied our approach to three public expression analyses, and found disease-associated gene sets not previously discovered in these studies.

2 Methods

2.1 The data

For definiteness, we present our notation using the language of gene expression experiments. Let g , h , r , and s denote individual genes and G be a set of genes. The cardinality of G is denoted $|G|$, or sometimes p . That is the same letter we use for p -value, but the usages are distinct enough that there should be no confusion. Our experiment has n subjects. The subjects may represent patients, cell cultures, or tissue samples.

The expression level for gene g in subject i is X_{gi} , and Y_i is the target variable on subject i . Y_i is often a treatment, disease, or other phenotype. We center the variables so that

$$\sum_{i=1}^n Y_i = \sum_{i=1}^n X_{gi} = 0, \quad \forall g. \quad (1)$$

The X_{gi} are not necessarily raw expression values, nor are they restricted to microarray values. In addition to the centering (1) they could have been scaled to have a given mean square. The scaling factor for X_{gi} might even depend on the sample variance for some genes $h \neq g$ if we thought that shrinking the variance for gene j towards the others would yield a more stable test statistic (Smyth, 2005). We might equally use a quantile transformation, replacing the j 'th largest of the raw X_{gi} by $\Phi^{-1}((j - 1/2)/n)$ where Φ is the Gaussian cumulative distribution function. Further preprocessing may be advised to handle outliers in X or Y . We do require that the preprocessing of the X 's does not depend on the Y 's and vice versa.

2.2 Test statistics

Our measure of association for gene g on our treatment of interest is

$$\hat{\beta}_g = \frac{1}{n} \sum_{i=1}^n X_{gi} Y_i. \quad (2)$$

If both X_{gi} and Y_i are centered and standardized to have variance 1, then $\hat{\beta}_g = \hat{\rho}_g$, the sample correlation between Y and gene g . The usual t -statistic for testing a linear relationship between these variables is $t_g \equiv \sqrt{n-2} \hat{\rho}_g / (1 - \hat{\rho}_g^2)^{1/2}$, which is a monotone transformation of $\hat{\rho}_g$.

For reasons of power and interpretability, we apply gene set testing methods instead of just testing individual genes. Linear and quadratic test statistics have been found to be the best performers for gene set enrichment analyses; we thus consider two statistics for our approach:

$$\hat{T}_{G,w} = \sum_{g \in G} w_g \hat{\beta}_g \quad \text{and} \quad \hat{C}_{G,w} = \sum_{g \in G} w_g \hat{\beta}_g^2.$$

The statistic $\hat{T}_{G,w}$ can approximate the JG score of Jiang and Gentleman (2007). The JG score is $(1/\sqrt{|G|}) \sum_{g \in G} t_g$. Taking $w_g = \sqrt{n-2}/(\text{sd}(X_g)\text{sd}(Y))$, where sd denotes a standard deviation, weights genes similarly to the JG score. Although $\hat{T}_{G,w}$ with these weights sums statistics equivalent to t statistics, it is not exactly equivalent to the sum of those statistics because of the way $\hat{\rho}_g$ appears in the denominator of each t_g .

The statistic $\hat{C}_{G,w}$ is a weighted sum of squared sample covariances. Ackermann and Strimmer (2009) conducted an extensive simulation of gene set methods and found good results for quadratic combinations of per gene test statistics.

The letters T and C are mnemonics for the t and χ^2 distributions that resemble the permutation distributions of these quantities. The w_g are scalar weights. For the quadratic statistics we will suppose that $w_g \geq 0$. We won't need that condition to find moments of $C_{G,w}$, but because we will compare $C_{G,w}$ to a χ^2 distribution, it is reasonable to avoid negative weights. Non-negative weights are also used to simplify our algorithm.

Although linear and quadratic test statistics are fairly restricted, they do allow a reasonable amount of customization through the weights w_g , and they are very interpretable compared to more ad hoc statistics.

2.3 Permutation procedure

A permutation of $\{1, 2, \dots, n\}$ is a reordering of $\{1, 2, \dots, n\}$. There are $n!$ permutations. We call π a *uniform random permutation* of $\{1, 2, \dots, n\}$ if it equals each distinct permutation with probability $1/n!$.

In a permutation analysis, we replace Y_i by \tilde{Y}_i where $\tilde{Y}_i = Y_{\pi(i)}$ for $i = 1, \dots, n$. Then $\tilde{\beta}_g = (1/n) \sum_{i=1}^n X_{gi} \tilde{Y}_i$, and when \tilde{Y} is substituted for Y , $\hat{T}_{G,w}$ becomes $\tilde{T}_{G,w}$ and $\hat{C}_{G,w}$ becomes $\tilde{C}_{G,w}$.

The $n!$ different permutations form a reference distribution from which we can compute p -values. There are often so many possible permutations that we cannot calculate or use all of them. Instead, we independently sample uniform random permutations M times, getting statistics $\tilde{C}_m = \tilde{C}_{G,w,m}$, and similarly \tilde{T}_m , for $m = 1, \dots, M$. We then compute p -values by comparing our observed statistics to our permutation distribution:

$$\begin{aligned} p_Q &= \frac{\#\{\tilde{C}_m \geq \hat{C}\} + 1}{M + 1} & p_C &= \frac{\#\{|\tilde{T}_m| \geq |\hat{T}|\} + 1}{M + 1} \\ p_L &= \frac{\#\{\tilde{T}_m \leq \hat{T}\} + 1}{M + 1}, \quad \text{or} & p_R &= \frac{\#\{\tilde{T}_m \geq \hat{T}\} + 1}{M + 1}, \end{aligned}$$

where p_Q and p_C are p -values for two-sided inferences on the quadratic and linear statistic, respectively, and p_L (left) and p_R (right) are for one-sided inferences based on the linear statistic. We use the mnemonic C in p_C to denote the central (or two-sided) p -value, which corresponds to a central confidence interval. The $+1$ in numerator and denominator of the p -values corresponds to counting the sample test statistic as one of the permutations. That is, we automatically include an identity permutation.

2.4 Permutation disadvantages

There are three main disadvantages to permutation-based analyses: cost, randomness, and granularity.

Testing many sets of genes becomes computationally expensive for two reasons. First, there are many test statistics to calculate in each permuted version of the data. Second, to allow for multiplicity adjustment, we require small nominal p -values to draw inference about our sets, which in turn requires a large number of permutations. That is, to obtain a small adjusted p -value (e.g., via FDR, FWER, Bonferroni methods), one first needs a small enough raw p -value. In order to obtain small raw p -values, the number of permutations (M) must be large, thereby increasing computational cost.

Because permutations are based on a random shuffling of the data, there is a chance that we will obtain a different p -value for our set of interest each time we run our permutation analysis. That is, our inference is subject to a given random seed.

Permutations also have a granularity problem. If we do M permutations, then the smallest possible p -value we can attain is $1/(M + 1)$. At or below this minimum p -value permutation tests have no power. Knijnenburg et al. (2009) suggest that for a reliable p -value, there should be at least 10 permuted values more extreme than the sample. That requires $M \approx 10/p$ and when it is necessary, due to test multiplicity, to use small p such as 10^{-6} or smaller, the permutation approach becomes computationally expensive. We call this the *sample granularity* problem.

There is also a *population granularity* problem. In an experiment with n observations, the smallest possible p -value is at least $1/n!$. Sometimes the at-

tainable minimum is much larger. For instance, when the target variable Y is binary with $n/2$ positive and $n/2$ negative values then the smallest possible p -value is $1/\binom{n}{n/2}$. For $n = 10$ we necessarily have $p \geq 1/252$. Rotation sampling methods such as ROAST are able to get around this population granularity problem (Wu et al., 2010). Increased Monte Carlo sampling can mitigate the sample granularity problem but not the population granularity problem.

Another aspect of the granularity problem is that permutations give us no basis to distinguish between two gene sets that both have the same p -value $1/(M + 1)$. There may be many such gene sets, and they have meaningfully different effect sizes. Many current approaches solve this problem by ranking significantly enriched gene sets by their corresponding test statistics. This practice only works if all test statistics have the same null distribution and correlation structure, which is not the case for many current GSEA methods. Additionally, the resulting broken ties do not have a p -value interpretation and cannot be directly used in multiple testing methods. To break ties in this way also requires the retention of both a p -value and a test statistic for inference, rather than just one value.

Because of each of these limitations of permutation testing, there is a need to move beyond permutation-based GSEA methods. The methods we present below are not as computationally expensive, random, or granular as their permutation counterparts. Our proposal results in a single number on the p -value scale.

2.5 Moment based reference distributions

To avoid the issues discussed above, we approximate the distribution of the permuted test statistics $\tilde{T}_{G,w}$ by Gaussians or by rescaled beta distributions. For quadratic statistics $\tilde{C}_{G,w}$ we use a distribution of the form $\sigma^2 \chi^2_{(\nu)}$ choosing σ^2 and ν to match the second and fourth moments of $\tilde{C}_{G,w}$ under permutation.

For the Gaussian treatment of $\tilde{T}_{G,w}$ we find $\sigma^2 = \text{var}(\tilde{T}_{G,w})$ under permutation using equation (5) of Section 3.3 and then report the p -value

$$p = \Pr(\mathcal{N}(0, \sigma^2) \leq \hat{T}_{G,w}),$$

where $\hat{T}_{G,w}$ is the observed value of the linear statistic. The above is a left tail p -value. Two-tailed and right-tailed p values are analogous.

When we want something sharper than the normal distribution, we can use a scaled Beta distribution, of the form $A + (B - A)\text{beta}(\alpha, \beta)$. The $\text{beta}(\alpha, \beta)$ distribution has a continuous density function on $0 < x < 1$ for $\alpha, \beta > 0$. We choose A , B , α and β by matching the upper and lower limits of $\tilde{T}_{G,w}$, as well

as its mean and variance. Using equation (5) from our theory section we have

$$\begin{aligned}
A &= \min_{\pi} \frac{1}{n} \sum_{i=1}^n \sum_{g \in G} w_g X_{gi} Y_{\pi(i)}, \\
B &= \max_{\pi} \frac{1}{n} \sum_{i=1}^n \sum_{g \in G} w_g X_{gi} Y_{\pi(i)}, \\
\alpha &= \frac{A}{B-A} \left(\frac{AB}{\text{var}(\tilde{T}_{G,w})} + 1 \right), \quad \text{and} \\
\beta &= \frac{-B}{B-A} \left(\frac{AB}{\text{var}(\tilde{T}_{G,w})} + 1 \right).
\end{aligned} \tag{3}$$

The observed left-tailed p -value is

$$p = \Pr\left(\text{beta}(\alpha, \beta) \leq \frac{\hat{T}_{G,w} - A}{B - A}\right).$$

It is easy to find the permutations that maximize and minimize $\tilde{T}_{G,w}$ by sorting the X and Y values appropriately as described in Section 3.3. The result has $A < 0 < B$. For the beta distribution to have valid parameters we must have $\sigma^2 < -AB$. From the inequality of Bhatia and Davis (2000), we know that $\sigma^2 \leq -AB$. There are in fact degenerate cases with $\sigma^2 = -AB$, but in these cases $\tilde{T}_{G,w}$ only takes one or two distinct values under permutation, and those cases are not of practical interest.

Like us, Zhou et al. (2009) have used a beta distribution to approximate a permutation. They used the first 4 moments of a Pearson curve for their approach. Fitting by moments in the Pearson family, it is possible to get a beta distribution whose support set (A, B) does not even include the observed value. That is, the observed value is even more extreme than it would have to be to get $p = 0$; it is almost like getting $p < 0$. We chose (A, B) based on the upper and lower limits of $\tilde{T}_{G,w}$ to prevent our observed test statistic from falling outside the range of possible values of our reference distribution (Section 3.3).

For the quadratic test statistic $\hat{C}_{G,w}$ we use a $\sigma^2 \chi_{(\nu)}^2$ reference distribution reporting the two-tailed p -value $\Pr(\sigma^2 \chi_{(\nu)}^2 \geq \hat{C}_{G,w})$ after matching the first and second moments of $\sigma^2 \chi_{(\nu)}^2$ to $\mathbb{E}(\tilde{C}_{G,w})$ and $\mathbb{E}(\tilde{C}_{G,w}^2)$ respectively. The parameter values are

$$\nu = 2 \frac{\mathbb{E}(\tilde{C}_{G,w})^2}{\text{var}(\tilde{C}_{G,w})} \quad \text{and} \quad \sigma^2 = \frac{\mathbb{E}(\tilde{C}_{G,w})}{\nu} = \frac{\text{var}(\tilde{C}_{G,w})}{2\mathbb{E}(\tilde{C}_{G,w})}.$$

Our formulas for $\mathbb{E}(\tilde{C}_{G,w})$ and $\mathbb{E}(\tilde{C}_{G,w}^2)$ under permutation are given in equation (4) of Section 3.1. Those formulas use $\mathbb{E}(\tilde{\beta}_g^2)$ and $\text{cov}(\tilde{\beta}_g^2, \tilde{\beta}_h^2)$ which we give in Corollaries 1 and 2 of Section 3.1.

All of our reference distributions are continuous and unbounded and hence they avoid the granularity problem of permutation testing. We have prepared

a publicly available Bioconductor (Gentleman et al., 2004) package, npGSEA, which implements our algorithm and calculates the corresponding statistics discussed in this section.

3 Theoretical results

3.1 Permutation moments of test statistics

Under permutation, $\mathbb{E}(\tilde{Y}_i) = 0$ by symmetry, and so $\mathbb{E}(\tilde{\beta}_g) = 0$ too. We easily find that,

$$\begin{aligned}\mathbb{E}(\tilde{T}_{G,w}) &= 0, \\ \text{var}(\tilde{T}_{G,w}) &= \sum_{g \in G} \sum_{h \in G} w_g w_h \text{cov}(\tilde{\beta}_g, \tilde{\beta}_h) \\ \mathbb{E}(\tilde{C}_{G,w}) &= \sum_{g \in G} w_g \mathbb{E}(\tilde{\beta}_g^2), \quad \text{and} \\ \text{var}(\tilde{C}_{G,w}) &= \sum_{g \in G} \sum_{h \in G} w_g w_h \text{cov}(\tilde{\beta}_g^2, \tilde{\beta}_h^2).\end{aligned}\tag{4}$$

The means, variances and covariances in (4) are taken with respect to the random permutations with the data X and Y held fixed. We adopt the convention that moments of permuted quantities are taken with respect to the permutation and are conditional on the X 's and Y 's. This avoids cumbersome expressions like $\mathbb{E}(\tilde{\beta}_g^2 \mid X_{gi}, Y_i, g \in G)$.

We will need the following even moments of X and Y :

$$\begin{aligned}\mu_2 &= \frac{1}{n} \sum_{i=1}^n Y_i^2, \quad \mu_4 = \frac{1}{n} \sum_{i=1}^n Y_i^4, \\ \bar{X}_{gh} &= \frac{1}{n} \sum_{i=1}^n X_{gi} X_{hi}, \quad \text{and} \\ \bar{X}_{ghrs} &= \frac{1}{n} \sum_{i=1}^n X_{gi} X_{hi} X_{ri} X_{si}\end{aligned}$$

for $g, h, r, s \in G$. Although our derivations involve $O(p^4)$ different moments when the gene set G has p genes, our computations do not require all of those moments.

Lemma 1. *For an experiment with $n \geq 2$ including genes g and h ,*

$$\mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h) = \frac{\mu_2 \bar{X}_{gh}}{n-1}.$$

Proof. See Appendix 1. \square

\square

Corollary 1. For an experiment with $n \geq 2$ including genes g and h ,

$$\text{cov}(\tilde{\beta}_g, \tilde{\beta}_h) = \mu_2 \bar{X}_{gh} / (n-1).$$

Proof. This follows from Lemma 1 because $\mathbb{E}(\tilde{\beta}_g) = 0$. \square

From Corollary 1, we see that the correlation between permuted test statistics $\tilde{\beta}_g$ and $\tilde{\beta}_h$ is simply the correlation between expression values for genes g and h .

Lemma 2. For an experiment with $n \geq 4$ including genes g, h, r, s ,

$$\mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h \tilde{\beta}_r \tilde{\beta}_s) = \begin{pmatrix} \mu_2^2 \\ \mu_4 \end{pmatrix}^\top A^\top B \begin{pmatrix} \bar{X}_{ghrs}^* / n^2 \\ \bar{X}_{ghrs} / n^3 \end{pmatrix}$$

where $\bar{X}_{ghrs}^* = \bar{X}_{gh}\bar{X}_{rs} + \bar{X}_{gs}\bar{X}_{hr} + \bar{X}_{gr}\bar{X}_{hs}$, with A^\top given by

$$\begin{pmatrix} 0 & 0 & \frac{n}{n-1} & \frac{-n}{(n-1)(n-2)} & \frac{3n}{(n-1)(n-2)(n-3)} \\ 1 & \frac{-1}{n-1} & \frac{-1}{n-1} & \frac{2}{(n-1)(n-2)} & \frac{-6}{(n-1)(n-2)(n-3)} \end{pmatrix},$$

and

$$B = \begin{pmatrix} 0 & 1 \\ 0 & -4 \\ 1 & -3 \\ -2 & 12 \\ 1 & -6 \end{pmatrix}.$$

Proof. See Appendix 2. \square

The expression is complicated, but it is simple to compute; we need only two moments of Y , two cross-moments of X , and the 2×2 matrix $A^\top B$. The matrix A depends on the experiment through n . Using Lemma 2 we can obtain the covariance between $\tilde{\beta}_g^2$ and $\tilde{\beta}_h^2$.

Corollary 2. For an experiment with $n \geq 4$ and genes g, h ,

$$\text{cov}(\tilde{\beta}_g^2, \tilde{\beta}_h^2) = \begin{pmatrix} \mu_2^2 \\ \mu_4 \end{pmatrix}^\top A^\top B \begin{pmatrix} \bar{X}_{gghh}^* / n^2 \\ \bar{X}_{gghh} / n^3 \end{pmatrix} - \frac{\mu_2^2}{(n-1)^2} \bar{X}_{gg} \bar{X}_{hh},$$

where $\bar{X}_{gghh}^* = \bar{X}_{gg}\bar{X}_{hh} + 2\bar{X}_{gh}^2$ with A and B as given in Lemma 2.

Proof. The covariance is $\mathbb{E}(\tilde{\beta}_g^2 \tilde{\beta}_h^2) - \mathbb{E}(\tilde{\beta}_g^2) \mathbb{E}(\tilde{\beta}_h^2)$. Applying Lemma 2 to the first expectation and Lemma 1 to the other two yields the result. \square \square

3.2 Rotation moments of test statistics

Rotation sampling (Wedderburn, 1975; Langsrud, 2005) provides an alternative to permutations, and is justified if either X or Y has a Gaussian distribution. It is simplest to describe when $Y \sim \mathcal{N}(\mu, \sigma^2 I_n)$ and even simpler for $Y \sim \mathcal{N}(0, \sigma^2 I_n)$. In the latter case we can replace Y by $\tilde{Y} = QY$ where $Q \in \mathbb{R}^{n \times n}$ is a random orthogonal matrix (independent of both X and Y), and the distribution of our test statistics is unchanged under the null hypothesis that X and Y are independent.

Rotation tests work by repeatedly sampling from the uniform distribution on random orthogonal matrices and recomputing the test statistics using \tilde{Y} instead of Y . They suffer from sample granularity but not population granularity because Q has a continuous distribution (for $n \geq 2$).

To take account of centering we need to use a rotation test appropriate for $Y \sim \mathcal{N}(\mu, \sigma^2 I_n)$. Langsrud (2005) does this by choosing rotation matrices that leave the population mean fixed. He rotates the data in an $n - 1$ dimensional space orthogonal to the vector 1_n . To get such a rotation matrix, he first selects an orthogonal contrast matrix $W \in \mathbb{R}^{n \times (n-1)}$. This matrix satisfies $W^T W = I_{n-1}$ and $W^T 1_n = 0_{n-1}$. Then he generates a uniform random rotation $Q^* \in \mathbb{R}^{(n-1) \times (n-1)}$ and delivers $\tilde{Y} = QY$, where $Q = \frac{1}{n} 1_n 1_n^T + W Q^* W^T$. More generally if $Y \sim \mathcal{N}(Z\gamma, \sigma^2 I_n)$, for a linear model $Z\gamma$, Langsrud (2005) shows how to rotate Y in the residual space of this model, leaving the fits unchanged.

Wu et al. (2010) have implemented rotation sampling for microarray experiments in their method, ROAST. They speed up the sampling by generating a random vector instead of a random matrix. For some tests, permutations and rotations have the same moments, and so our approximations are approximations of rotation tests as much as of permutation tests.

Our rotation method approximation performs very similarly to the permutation method. We let $\tilde{Y} = QY$ for $Q = (\frac{1}{n} 1_n 1_n^T + W Q^* W^T)$ where Q^* is a uniform random $n - 1 \times n - 1$ rotation matrix and the contrast matrix $W \in \mathbb{R}^{n \times (n-1)}$ satisfies $W^T 1_n = 0_{n-1}$ and $W^T W = I_{n-1}$ and then $\tilde{\beta}$, \tilde{T} and \tilde{C} are defined as for permutations, substituting \tilde{Y} for Y .

The variance of the quadratic test statistic depends on *which* contrast matrix W one chooses, and it cannot always match the permutation variance. This difference disappears asymptotically as $n \rightarrow \infty$.

Lemma 3. *For an experiment with $n \geq 2$ including genes g and h , the moments $\mathbb{E}(\tilde{\beta}_g)$ and $\mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h)$ are identical to their permutation counterparts, regardless of the choice for W .*

Proof. See Appendix 3 and 4. \square

Corollary 3. *For an experiment with $n \geq 2$, $\mathbb{E}(\tilde{T}_{G,w})$, $\text{var}(\tilde{T}_{G,w})$ and $\mathbb{E}(\tilde{C}_{G,w})$ are the same whether \tilde{Y} is formed by permutation or rotation of Y .*

3.3 Computation and costs

To facilitate computation for the linear statistic, we reduce each gene set to a single pseudo-gene $X_{Gi} = \sum_{g \in G} w_g X_{gi}$ and then let

$$\bar{X}_G = \frac{1}{n} \sum_{i=1}^n X_{Gi} \quad \text{and} \quad \bar{X}_{GG} = \frac{1}{n} \sum_{i=1}^n X_{Gi}^2.$$

The weights w have been absorbed into the pseudo-gene to simplify notation. We define

$$\begin{aligned} \hat{\beta}_G &= \sum_{g \in G} w_g \hat{\beta}_g = \frac{1}{n} \sum_i X_{Gi} Y_i, \quad \text{and} \\ \tilde{\beta}_G &= \sum_{g \in G} w_g \tilde{\beta}_g = \frac{1}{n} \sum_i X_{Gi} \tilde{Y}_i. \end{aligned}$$

Our permuted linear test statistic is $\tilde{T}_{G,w} = \tilde{\beta}_G$, with

$$\text{var}(\tilde{T}_{G,w}) = \text{var}(\tilde{\beta}_G) = \frac{\mu_2}{n-1} \bar{X}_{GG}. \quad (5)$$

For the beta approximation, we need the range of $\tilde{T}_{G,w}$. Let the sorted Y values be $Y_{(1)} \leq Y_{(2)} \leq \dots \leq Y_{(n)}$ and the sorted X_{Gi} values be $X_{G(1)} \leq X_{G(2)} \leq \dots \leq X_{G(n)}$. Then the range of $\tilde{T}_{G,w}$ is $[A, B]$, where

$$A = \frac{1}{n} \sum_{i=1}^n X_{G(i)} Y_{(n+1-i)}, \quad \text{and} \quad B = \frac{1}{n} \sum_{i=1}^n X_{G(i)} Y_{(i)}.$$

For a $\sigma t_{(\nu)}$ reference distribution we would also need $\mathbb{E}(\tilde{T}_{G,w}^4) = \mathbb{E}(\tilde{\beta}_G^4)$. We can apply Lemma 2 to the pseudo-gene resulting in

$$\mathbb{E}(\tilde{\beta}_G^4) = \begin{pmatrix} \mu_2^2 \\ \mu_4 \end{pmatrix} A^\top B \begin{pmatrix} 3\bar{X}_{GG}^2/n^2 \\ \bar{X}_{GGGG}/n^3 \end{pmatrix}, \quad (6)$$

where $\bar{X}_{GGGG} = \frac{1}{n} \sum_{i=1}^n X_{Gi}^4$.

We considered using a $\sigma t_{(\nu)}$ reference distribution for $\tilde{T}_{G,w}$, taking into account the fourth moment of $\tilde{T}_{G,w}$ (6). We have often (in fact usually) found that $\mathbb{E}(\tilde{T}_{G,w}^4) < 3\mathbb{E}(\tilde{T}_{G,w}^2)^2$; that is, lighter tails than the normal. This implies a negative kurtosis for the permutation distribution, and t distributions have positive kurtosis. For this reason we use a beta approximation and not a t approximation.

For the quadratic statistic we have found it useful to replace X_{gi} by $\sqrt{w_g} X_{gi}$ in precomputation. That step is only valid for non-negative w_g , but those are the ones of most interest. Then we use formulas for $\mathbb{E}(\tilde{C}_{G,w})$ and $\text{var}(\tilde{C}_{G,w})$ with all $w_g = w_h = 1$ (4).

Now we consider the computational cost. The cost to compute all of the X_{Gi} is dominated by np multiplications. It then takes n more multiplications to get $\hat{\beta}_G$ and another n to get \tilde{X}_{GGe} . It costs n multiplications to get all the μ_2 , except that step done once can be used for all gene sets. The cost for the Gaussian approximation $\mathcal{N}(0, \text{var}(\tilde{T}_{G,w}))$ is dominated by $n(p+2)$ multiplications.

For the beta approximation there is also a cost proportional to $n \log(n)$ in the sorting to compute limits A and B . That adds a cost comparable to a multiple of $\log(n)$ permutations. We judge that the cost of sorting is usually minor for n and p of interest in bioinformatics.

A permutation analysis requires nM multiplications, after computing X_{Gi} , for a total of $n(M+p)$. It is very common for p to be a few tens and M to be many thousands or more. Then we can simplify the costs to $n(M+p) \approx nM$ and $n(2+p) \approx np$. The moment method costs about as much as doing p permutations. When the gene set has tens of genes and the permutation method uses many thousands or even several million permutations, the computational cost is quite large.

The pseudo-gene technique is more expensive for the quadratic statistics. The dominant cost in computing $\hat{C}_{G,w}$ is still the np multiplications required to compute $\hat{\beta}_g$ for $g \in G$. We can also compute $\mathbb{E}(\tilde{C}_{G,w})$ in about this amount of work.

The cost of computing $\text{var}(\tilde{C}_{G,w})$ by a straightforward algorithm is at least np^2 , because we need \tilde{X}_{gh} and \tilde{X}_{gghh} for all $g, h \in G$. Some parts of that computation can be sped up to $O(np)$ by rewriting the expression as described in Appendix 5. One of the terms however does not reduce to $O(np)$. A straightforward implementation costs $O(np^2)$ while an alternative expression costs $O(n^2p)$. The latter is valuable in settings where the gene sets are large compared to the sample size. In the former case, the moment approximation has cost comparable to $O(p^2)$ permutations. If $n < p$ then the latter case is like np permutations, so the quadratic cost is comparable to on the order of $p * \min(n, p)$ permutations.

We have verified our moment formulas by finding that they match values found by enumerating all $n!$ permutations, for some simulated data sets with small n . During testing, we also compared permutation and our approximate p -values on simulated data. We saw a close match but think that an illustration on real data is more compelling. Section 4 makes comparisons using three genome wide expression studies in Parkinson’s Disease (PD) patients: Moran et al. (2006), Scherzer et al. (2007) and Zhang et al. (2005).

4 Parkinson’s Disease

We illustrate our method using publicly available data from three expression studies in Parkinson’s Disease (PD) patients (Moran *et al.*, 2006, Zhang *et al.*, 2005, and Scherzer *et al.*, 2007; Table 1). All three experiments contain genome wide expression values measured via a microarray experiment. PD is a common neurodegenerative disease; clinical symptoms often include rigidity, resting tremor and gait instability (Abou-Sleiman et al., 2006). Pathologically, PD

Reference	Tissue	# Affected	# Controls
Moran	Substantia nigra	29	14
Zhang	Substantia nigra	18	11
Scherzer	Blood	47	21

Table 1: Three data sets used for non-permutation GSEA

is characterized by neuronal-loss in the substantia nigra and the presence of α -synuclein protein aggregates in neurons (Abou-Sleiman et al., 2006).

Using a selected set from the Broad Institute’s mSigDB v3.1 (Subramanian et al., 2005) and the presence of PD as a response variable from the Zhang et al. (2005) dataset, we visualized both permutation distributions and our approximation of these distributions (Figure 1). As discussed above, we use a linear test statistic, $\hat{T}_{G,w} = \sum_{g \in G} \hat{\beta}_g$, and a quadratic test statistic, $\hat{C}_{G,w} = \sum_{g \in G} \hat{\beta}_g^2$, where $\hat{\beta}_g$ is a sample covariance between gene expression and, in this case, disease status. Figure 1 shows these two test statistics with a histogram of 99,999 recomputations of those statistics for permutations of treatment status versus gene expression. In principle, histograms of permuted test statistics can be very complicated, but in practice, they often resemble familiar parametric distributions, as in Figure 1.

Using the fitted normal distribution to determine the rarity of the observed gene set statistic results in a two-tailed p -value of 0.0604 for the linear statistic while permutations yield $p = 0.0595$. A fitted $\sigma^2 \chi^2_{(\nu)}$ distribution results in $p = 0.0425$ for the sum of squares gene set statistic, while permutations yield $p = 0.0458$. The p -values are quite close despite the somewhat higher peak for the permutation histogram relative to the χ^2 density.

We compared our non-permutation p -values to p -values for linear and quadratic statistics for the 6,303 gene sets from mSigDB’s curated gene sets and Gene Ontology (GO, Ashburner *et al.*, 2000) gene sets collections (v3.1). One gene set was removed because it contained only one gene in our experiments. The average size of these gene sets is 79.40 genes. For our gold standard we ran 999,999 permutations of the linear statistic and 499,999 permutations of the quadratic statistic. For all of our permutations, we first calculated the observed test statistic for each of the 6,303 gene sets and then permuted the Y_i ’s M times to obtain $6,303 \times M$ permuted test statistics. We next compared the pre-computed test statistic vector to our matrix of permuted test statistics.

For each set, we computed left-sided p -values, p_L , for the linear statistic and two-sided p -values, p_Q , for the quadratic statistic using these permutations. We also computed the normal and beta approximations of p_L with our method. (Figure 2, left panel). We converted these one-sided p -values to two-sided p -values via $p = 2 \min(p_L, 1 - p_L)$. The beta approximation p -values are almost identical to the permutation p -values.

For our quadratic test statistic, we fit our moment based $\sigma^2 \chi^2_{(\nu)}$ approximation and computed two-sided tailed p -values across all sets (Figure 2, right

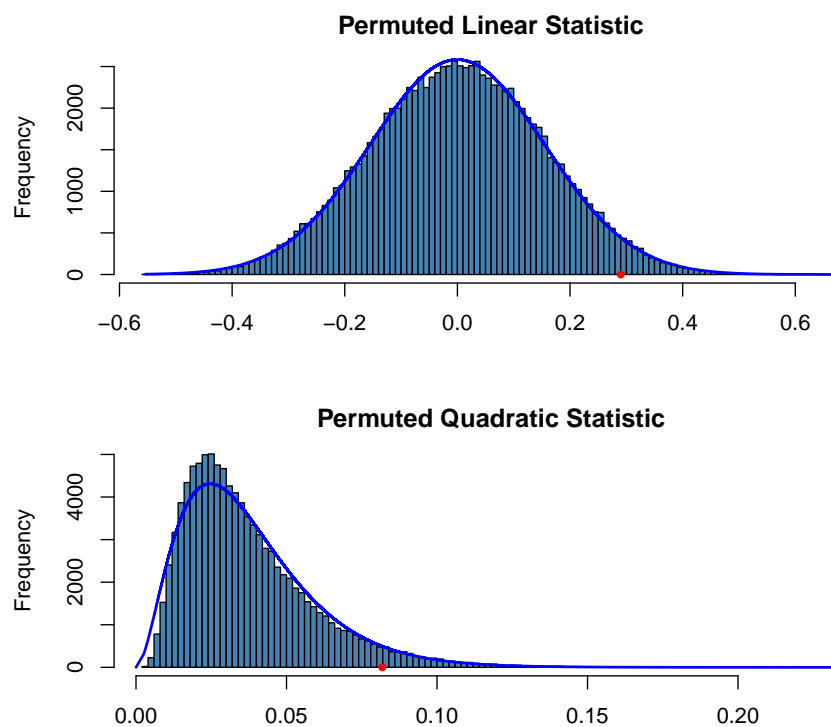


Figure 1: Top panel shows a permutation histogram for a linear test statistic for the steroid hormone signaling pathway gene set as described in the text. The bottom panel shows a quadratic test statistic. Solid red dots indicate the observed values and curves indicate parametric fits, based on normal and χ^2 distributions.

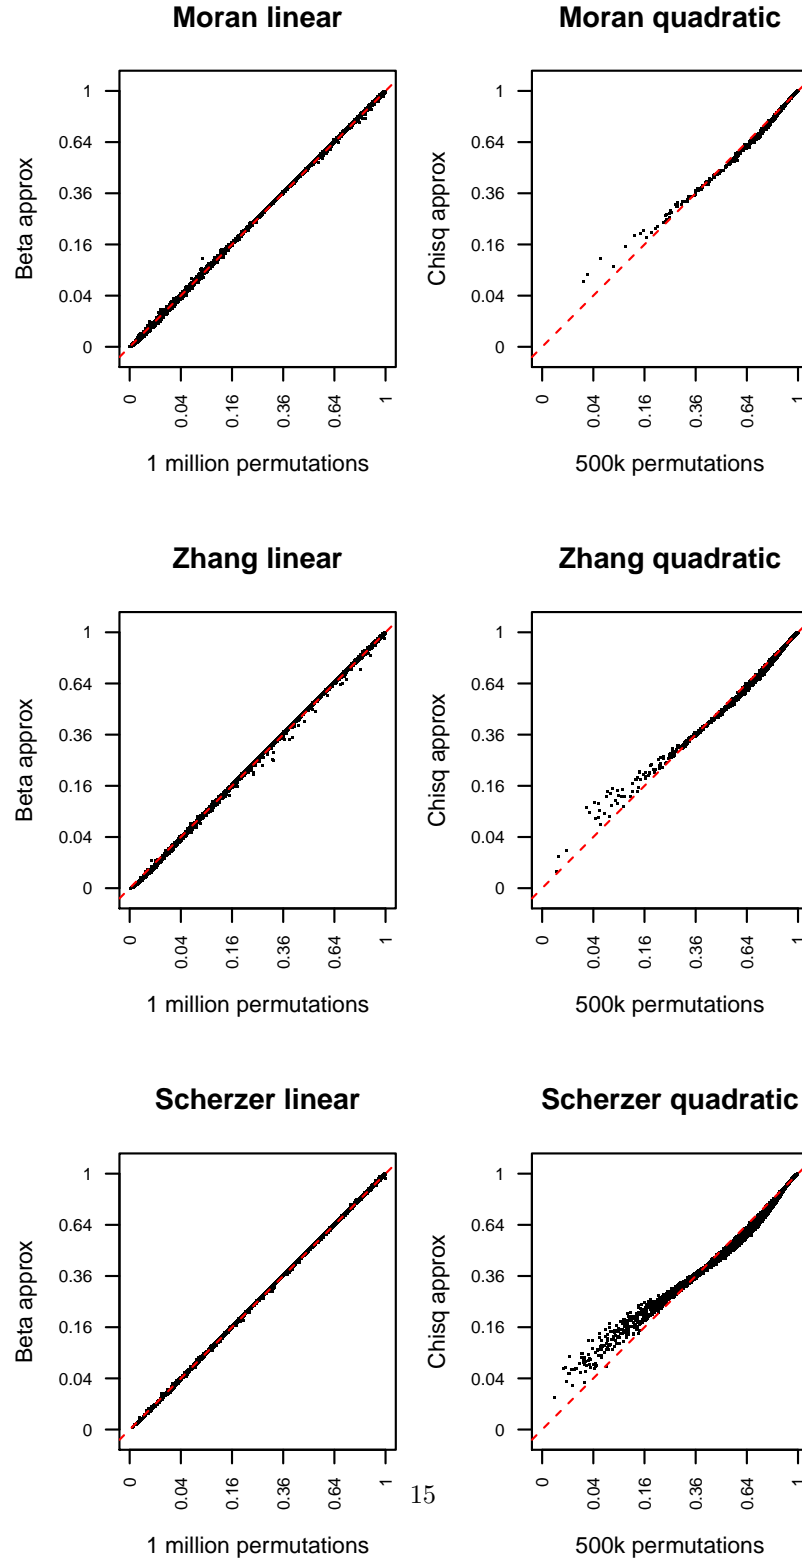


Figure 2: Permutation p -values (x-axis) versus moment-based p -values (y-axis) for 6,303 gene sets. The left column represents results for a linear test statistic, the right column for sum of squares. Data come from three genome-wide expression studies. We applied the non-linear transformation $p^{1/2}$ to stretch the

Reference	Normal p_L	Beta p_L	Normal p_C	Beta p_C	Chisq p_Q
Moran	0.99991	0.99997	0.99973	0.99991	0.978
Zhang	0.99996	0.99997	0.99983	0.99991	0.990
Scherzer	0.99998	0.99999	0.99991	0.99997	0.994

Table 2: Spearman correlations between gold standard (999,999 and 499,999 permutations for linear and quadratic statistics) and approximation p -values. p_L and p_C represent results for one and two-tailed linear test statistics, respectively. Chisq p_Q represents results for the sum of squares analysis.

panel). We see that the smallest χ^2 non-permutation p -values are slightly conservative. This may reflect the boundedness of the permutation distribution combined with the unbounded right tail of the χ^2 distribution.

In each of the three experiments, there is a tight correlation between the permutation-based p -values of all sets and both of our moment-based methods (Table 2). The beta and normal approximations are almost identical. Our beta approximations are slightly closer to the gold standard than the normal approximations, but not by a practically important amount. The beta approximation has shorter tails than the Gaussian approximation. It yielded p -values somewhat smaller than permutations did, while the Gaussian approximation yielded p -values somewhat larger than the permutations did. The χ^2 approximations also reproduce the ranking of the gold standard quite well, though not as well as the normal and beta approximations to the linear statistic.

For these data sets and 6,303 gene sets, both of the linear statistics, which have more or less the same rank-ordering of p -values as 999,999 permutations, could be approximated in about the amount of time it takes to compute 100 permutations (Table 3, top block). Our gene sets had an average size of about 80 genes. This lead us to expect that the cost of the linear approximation would be comparable to doing 80 permutations. We found that the Gaussian approximation cost about as much as 100 permutations. While this is a close match, we remark that the time to do M permutations is nearly an affine function $a + bM$ with positive intercept a . At such small M the overhead costs dominated the total cost making the per permutation costs hard to resolve. The beta approximation was slightly slower than the Gaussian one because it involves the sorting of the data.

The χ^2 approximation to the quadratic statistic has a computational cost about as much as 35,000 to 45,000 permutations, yet has a similar rank-ordering of p -values 499,999 permutations (Table 3, bottom block). For the quadratic statistic we expected our algorithm to cost as much as doing a number of permutations equal to a small multiple of the mean square gene set size. It cost about as much as 35,000 to 45,000 permutations while the mean square set size was 27,171.

After applying our permutation approximation methods to each dataset in 6,303 mSigDB gene sets, we found many significantly enriched gene sets, even

Method	Moran	Zhang	Scherzer
$M = 100$	31.03	29.84	34.71
$M = 500$	31.95	32.49	35.54
$M = 1,000,000$	5010.17	4434.77	3933.15
Normal	29.74	27.00	34.66
Beta	30.79	31.88	37.89
$M = 30,000$	9146.27	7217.59	11808.02
$M = 40,000$	12256.54	9636.06	16545.60
$M = 50,000$	16833.08	12564.06	21480.80
$M = 500,000$	149588.37	129667.73	187067.91
χ^2	11020.62	10600.82	12677.15

Table 3: Time in seconds for p -value calculations for 6,303 gene sets in three genome-wide expression studies. Linear statistic results with $M = 100$, $M = 500$, and $M = 1,000,000$ permutations, and the normal and beta approximations are in the top block. Timings for the quadratic statistic with $M = 30,000$, $M = 40,000$, $M = 50,000$, and $M = 500,000$ permutations, and the χ^2 approximation are presented in the bottom block.

after correcting for multiple testing (two-sided adjusted p -value < 0.05). The most significantly enriched sets are associated with metabolism and mitochondrial function, neuronal transmitters and serotonin, epigenetic modifications, and the transcription factor FOXP3 Supplemental Table 1¹ Each of these categories has some previously discovered association with PD, although not through traditional gene set methods (metabolism and mitochondrial function: Abou-Sleiman et al. (2006); neuronal transmitters and serotonin: Fox et al. (2009); epigenetic modifications: Berthier and Pulido (2013); FOXP3: Stone et al. (2009)). Through our new gene set enrichment method, we discovered a relationship between the expression of these gene sets and PD.

5 Discussion

Gene set methods are able to pool weak single gene signals over a set of genes to get a stronger inference. These methods and their corresponding permutation-based inferences are a staple of high throughput methods in genomics. Because an experiment for this purpose may have a few to hundreds of microarrays or RNA-seq samples, permutation can be computationally costly, and yet still result in granular p -values. In this paper, we introduce an approximation gene set method, which performs as well as permutation methods, in a fraction of the computation time and which generates continuous p -values.

Permutation methods have some valuable properties that our approach does not share. Permutation based inferences give exact p -values. Our approxima-

¹ <http://statweb.stanford.edu/~owen/reports/SupplementalTable1.xls>

tions are not ordinarily exact because the permutation histogram is not in the parametric family we use.

The second advantage of permutations is that they apply to arbitrarily complicated statistics. In our view, many of those complicated statistics are much harder to interpret and are less intuitive than the plain sum and sum of squared statistics we present. Others have observed that simple linear and squared statistics outperform more complex approaches (Ackermann and Strimmer, 2009). Our method allows for the weighting of coefficients in our statistics, granting users access to additional useful and interpretable patterns.

Because of the disadvantages discussed above, there has long been interest in finding approximations to permutation tests. Eden and Yates (1933) noticed that the permutation distribution closely matched a parametric distribution that one would get running an F -test on the same data. It has also been known since the 1940s that the permutation distribution of the linear test is asymptotically normal as n increases (Good, 2004). More recently, Knijnenburg et al. (2009) approach the granularity issue by taking a random sample of permutations and fitting a generalized extreme value (GEV) distribution to the tail of their distribution.

Our work differs from these previous permutation approximation approaches. We use Gaussian or beta distributions for the linear statistic and a χ^2 distribution for the quadratic statistic. These choices never place the observed test statistic strictly outside the possible range of our reference distribution. In this way, we also avoid nonsensical p -values.

We have developed a new and intuitive method for gene set enrichment analysis that is computationally inexpensive, as accurate as permutation methods, and avoids the sample granularity issue. A Gaussian, beta, or χ^2 approximation gives a principled way to break ties among genes or gene sets whose test statistics are larger than any seen in the M permutations. We applied our moment based approximations to three human Parkinson’s Disease data sets and discovered the enrichment of several gene sets in this disease, none of which were mentioned in the original publications.

Acknowledgement

We thank Nicholas Lewin-Koh, Joshua Kaminker, Richard Bourgon, Sarah Kummerfeld, Thomas Sandmann, and John Robinson for helpful comments. ABO thanks Robert Gentleman, Jennifer Kesler and other members of the Bioinformatics and Computational Biology Department at Genentech for their hospitality during his sabbatical there.

Funding: JLL is funded by Genentech, Inc. ABO was supported by Genentech, Inc. and by Stanford University while on a sabbatical.

References

- Abou-Sleiman, P., Muqit, M., and Wood, N. (2006). Expanding insights of mitochondrial dysfunction in parkinsons disease. *Nat Rev Neurosci*, 7:207–219.
- Ackermann, M. and Strimmer, K. (2009). A general modular framework for gene set enrichment analysis. *BMC Bioinformatics*, 10:47–66.
- Anderson, T., Olkin, I., and Underhill, L. (1987). Generation of random orthogonal matrices. *SIAM Journal on Scientific and Statistical Computing*, 8(4):625–629.
- Berthier, A. Jimnez-Sinz, J. and Pulido, R. (2013). Pink1 regulates histone h3 trimethylation and gene expression by interaction with the polycomb protein eed/wait1. *Proc Natl Acad Sci USA*, 110(36):14729–34.
- Bhatia, R. and Davis, C. (2000). A better bound on the variance. *The American Mathematical Monthly*, 107(4):353–357.
- Eden, T. and Yates, F. (1933). On the validity of Fisher’s z -test when applied to an actual sample of non-normal values. *The Journal of Agricultural Science*, 23:6–7.
- Fox, S., Chuang, M., and Brotchie, J. (2009). Serotonin and parkinsons disease: On movement, mood, and madness. *Movement Disorders*, 24(9):1255–1266.
- Gentleman, R., Carey, V., Bates, D., Bolstad, B., Dettling, M., Dudoit, S., Ellis, B., Gautier, L., Ge, Y., Gentry, J., Hornik, K., Hothorn, T., Huber, W., Iacus, S., Irizarry, R., Leisch, F., Li, C., Maechler, M., Rossini, A., Sawitzki, G., Smith, C., Smyth, G., Tierney, L., Yang, J., and Zhang, J. (2004). Bioconductor: open software development for computational biology and bioinformatics. *Genome Biol*, 5(10):R80.1–R80.16.
- Goeman, J. J. and Bühlmann, P. (2007). Analyzing gene expression data in terms of gene sets: methodological issues. *Bioinformatics*, 23(8):980–987.
- Good, P. I. (2004). *Permutation, parametric, and bootstrap tests of hypotheses*. Springer, New York.
- Jiang, Z. and Gentleman, R. (2007). Extensions to gene set enrichment. *Bioinformatics*, 23(3):306–313.
- Knijnenburg, T. A., Wessels, L. F. A., Reinders, M. J. T., and Shmulevich, I. (2009). Fewer permutations, more accurate p-values. *Bioinformatics*, 25(12):i161–i168.
- Langsrud, O. (2005). Rotation tests. *Statistics and computing*, 15:53–60.
- Lehmann, E. L. and Romano, J. P. (2005). *Testing statistical hypotheses*. Springer.

- Mootha, V. K., Lindgren, C. M., Eriksson, K. F., Subramanian, A., Sihag, S., Lehar, J., Puigserver, P., Carlsson, E., Ridderstrale, M., Laurila, E., Houstis, N., Daly, M. J., Patterson, N., Mesirov, J. P., Golub, T. R., Tamayo, P., Spiegelman, B., Lander, E. S., Hirschhorn, J. N., Altshuler, D., and Groop, L. C. (2003). PGC-1 α -responsive genes involved in oxidative phosphorylation are coordinately downregulated in human diabetes. *Nature Genetics*, 34:267–273.
- Moran, L. B., Duke, D. C., Deprez, M., Dexter, D. T., Pearce, R. K. B., and Graeber, M. B. (2006). Whole genome expression profiling of the medial and lateral substantia nigra in Parkinsons disease. *Neurogenetics*, 7(1):1–11.
- Newton, M. A., Quintana, F. A., den Boon, J. A., Sengupta, S., and Ahlquist, P. (2007). Random-set methods identify distinct aspects of the enrichment signal in gene-set analysis. *The Annals of Applied Statistics*, pages 85–106.
- Owen, A. B. (2005). Variance of the number of false discoveries. *Journal of the Royal Statistical Society, Series B*, 67(3):411–426.
- Scherzer, C. R., AC, A. C. E., Morse, L. J., Liao, Z., Locascio, J. J., Fefer, D., Schwarzschild, M. A., Schlossmacher, M. G., Hauser, M. A., Vance, J. M., Sudarsky, L. R., Standaert, D. G., Growdon, J. H., Jensen, R. V., and Gullans, S. R. (2007). Molecular markers of early Parkinson’s disease based on gene expression in blood. *Proc Natl Acad Sci*, 104(3):955–60.
- Smyth, G. (2005). Limma: linear models for microarray data. In Gentleman, R., Carey, V., Dudoit, S., Irizarry, R., and Huber, W., editors, *Bioinformatics and Computational Biology Solutions Using R and Bioconductor*, pages 397–420. Springer, New York.
- Stone, D., Reynolds, A., Mosely, R., and Gendelman, H. (2009). Innate and adaptive immunity for the pathobiology of parkinson’s disease. *Antioxid Redox Signal*, 11(9):2151–2166.
- Subramanian, A., Tamayo, P., Mootha, V., Mukherjee, S., Ebert, B., Gillette, M., Paulovich, A., Pomeroy, S., Golub, T., Lander, E., and Mesirov, J. (2005). Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci USA*, 102(43):15545–50.
- Wedderburn, R. W. M. (1975). Random rotations and multivariate normal simulation. Technical report, Rothamsted Experimental Station.
- Wu, D., Lim, E., Vaillant, F., Asselin-Labat, M.-L., Visvader, J. E., and Smyth, G. K. (2010). Roast: rotation gene set tests for complex microarray experiments. *Bioinformatics*, 26(17):2176–2182.
- Zhang, Y., James, M., Middleton, F. A., and Davis, R. L. (2005). Transcriptional analysis of multiple brain regions in Parkinson’s disease supports the

involvement of specific protein processing, energy metabolism, and signaling pathways, and suggests novel disease mechanisms. *Am J Med Genet B Neuropsychiatr Genet*, 137B(1):5–16.

Zhou, C., Wang, H. J., and Wang, Y. M. (2009). Efficient moments-based permutation tests. *Advances in neural information processing systems*, 22:2277.

Appendix 1: Proof of Lemma 1

This appears in Owen (2005) but we prove it here to keep the paper self-contained. First

$$n^2 \mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h) = \sum_i \sum_{i'} X_{gi} X_{hi'} \mathbb{E}(\tilde{Y}_i \tilde{Y}_{i'})$$

Recall that $\mu_2 = \frac{1}{n} \sum_{i=1}^n Y_i^2$. Then

$$\mathbb{E}(\tilde{Y}_i \tilde{Y}_{i'}) = \begin{cases} \mu_2, & i' = i \\ -\frac{1}{n-1} \mu_2, & i' \neq i \end{cases}$$

and so

$$\begin{aligned} n^2 \mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h) &= \sum_i \sum_{i'} X_{gi} X_{hi'} \mathbb{E}(\tilde{Y}_i \tilde{Y}_{i'}) \\ &= \mu_2 \sum_i \sum_{i'} X_{gi} X_{hi'} \left(1_{i=i'} - \frac{1}{n-1} 1_{i \neq i'} \right) \\ &= \mu_2 \sum_i \sum_{i'} X_{gi} X_{hi'} \left(\frac{n}{n-1} 1_{i=i'} - \frac{1}{n-1} \right) \\ &= \frac{n}{n-1} \mu_2 \sum_i X_{gi} X_{hi} \\ &\equiv \frac{n^2}{n-1} \mu_2 \bar{X}_{gh}, \end{aligned}$$

proving Lemma 1. \square

Appendix 2: Proof of Lemma 2

The fourth moment contains terms of the form

$$X_{gi} X_{hj} X_{rk} X_{s\ell} \mathbb{E}(\tilde{Y}_i \tilde{Y}_j \tilde{Y}_k \tilde{Y}_\ell)$$

and there are different special cases depending on which pairs of indices among i, j, k and ℓ are equal. We need the following fourth moments of Y in which all

indices are distinct:

$$\begin{aligned}\mu_{4k} &= \mathbb{E}(\tilde{Y}_i^4) \\ \mu_{3k} &= \mathbb{E}(\tilde{Y}_i^3 \tilde{Y}_j) \\ \mu_{2p} &= \mathbb{E}(\tilde{Y}_i^2 \tilde{Y}_j^2) \\ \mu_{1p} &= \mathbb{E}(\tilde{Y}_i^2 \tilde{Y}_j \tilde{Y}_k) \\ \mu_{\emptyset} &= \mathbb{E}(\tilde{Y}_i \tilde{Y}_j \tilde{Y}_k \tilde{Y}_\ell),\end{aligned}$$

and where the subscripts are mnemonics for terms four of a kind, three of a kind, two pair, one pair and nothing special.

We can express all of these moments in terms of μ_2 and $\mu_4 = (1/n) \sum_{i=1}^n Y_i^4$. Each moment is a normalized sum over distinct indices. We can write these in terms of normalized sums over all indices. Many of those terms vanish because $\sum_i Y_i = 0$.

Let \sum^* represent summation over distinct indices, as in

$$\begin{aligned}\sum_{ij}^* f_{ij} &= \sum_{i=1}^n \sum_{j=1, j \neq i}^n f_{ij}, \\ \sum_{ijk}^* f_{ijk} &= \sum_{i=1}^n \sum_{j=1, j \neq i}^n \sum_{k=1, k \neq i, k \neq j}^n f_{ijk}\end{aligned}$$

and so on. We can write these sums in terms of unrestricted sums:

$$\begin{aligned}\sum_{ij}^* f_{ij} &= \sum_{ij} f_{ij} - \sum_i f_{ii} \\ \sum_{ijk}^* f_{ijk} &= \sum_{ijk} f_{ijk} - \sum_{ij} (f_{iij} + f_{iji} + f_{ijj}) + 2 \sum_i f_{iii}, \quad \text{and} \\ \sum_{ijk\ell}^* f_{ijk\ell} &= \sum_{ijk\ell} f_{ijk\ell} - \sum_{ijk} (f_{ijk i} + f_{ijk j} + f_{ijk k} + f_{ijk i} + f_{ijk j} + f_{ijk k}) \\ &\quad + \sum_{ij} (2(f_{ijjj} + f_{ijii} + f_{iiji} + f_{iiij}) + f_{ijij} + f_{ijji} + f_{iijj}) - 6 \sum_i f_{iiii}.\end{aligned}$$

See Gleich and Owen (2011) for details.

We will use the last expression in a context where $f_{ijk\ell}$ vanishes when summed over the entire range of any one of its indices. In that case

$$\sum_{ijk\ell}^* f_{ijk\ell} = \sum_{ij} (f_{ijij} + f_{ijji} + f_{iijj}) - 6 \sum_i f_{iiii}. \quad (7)$$

We also use the notation $n^{(k)} = n(n-1)(n-2) \cdots (n-k+1)$, often called ‘ n

to k factors', where k is a positive integer. Now

$$\begin{aligned}
\mu_{4k} &= \frac{1}{n} \sum_{i=1}^n Y_i^4 = \mu_4, \\
\mu_{3k} &= \frac{1}{n^{(2)}} \sum_{ij}^* Y_i^3 Y_j = \frac{1}{n^{(2)}} \left(\sum_{ij} Y_i^3 Y_j - \sum_i Y_i^4 \right) \\
&= -\frac{\mu_4}{n-1}, \\
\mu_{2p} &= \frac{1}{n^{(2)}} \sum_{ij}^* Y_i^2 Y_j^2 = \frac{1}{n^{(2)}} \left(\sum_{ij} Y_i^2 Y_j^2 - \sum_i Y_i^4 \right) \\
&= \frac{1}{n-1} (n\mu_2^2 - \mu_4), \quad \text{and} \\
\mu_{1p} &= \frac{1}{n^{(3)}} \sum_{ijk}^* Y_i^2 Y_j Y_k \\
&= \frac{1}{n^{(3)}} \left(\sum_{ijk} Y_i^2 Y_j Y_k - \sum_{ij} (2Y_i^3 Y_j + Y_i^2 Y_j^2) + 2 \sum_i Y_i^4 \right) \\
&= \frac{-n\mu_2^2 + 2\mu_4}{(n-1)(n-2)}.
\end{aligned}$$

Finally using (7), $n^{(4)}\mu_\emptyset$ equals

$$\sum_{ijkl}^* Y_i Y_j Y_k Y_\ell = 3 \sum_{ij} Y_i^2 Y_j^2 - 6 \sum_i Y_i^4 = 3n^2 \mu_2^2 - 6n\mu_4$$

so that

$$\mu_\emptyset = \frac{1}{(n-1)(n-2)(n-3)} (3n\mu_2^2 - 6\mu_4).$$

We may summarize these results via

$$\begin{pmatrix} \mu_{4k} \\ \mu_{3k} \\ \mu_{2p} \\ \mu_{1p} \\ \mu_\emptyset \end{pmatrix} = A \begin{pmatrix} \mu_2^2 \\ \mu_4 \end{pmatrix},$$

where the matrix A is given in the statement of Lemma 2.

Now

$$\begin{aligned}
n^4 \mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h \tilde{\beta}_r \tilde{\beta}_s) &= \sum_{ijk\ell} X_{gi} X_{hj} X_{rk} X_{s\ell} \mathbb{E}(\tilde{Y}_i \tilde{Y}_j \tilde{Y}_k \tilde{Y}_\ell) \\
&= \mu_{4k} \sum_i X_{gi} X_{hi} X_{ri} X_{si} \\
&\quad + \mu_{3k} \sum_{ij}^* (X_{gi} X_{hi} X_{ri} X_{sj} + X_{gi} X_{hi} X_{rj} X_{si} + X_{gi} X_{hj} X_{ri} X_{si} + X_{gj} X_{hi} X_{ri} X_{si}) \\
&\quad + \mu_{2p} \sum_{ij}^* (X_{gi} X_{hi} X_{rj} X_{sj} + X_{gi} X_{hj} X_{ri} X_{sj} + X_{gi} X_{hj} X_{rj} X_{si}) \\
&\quad + \mu_{1p} \sum_{ijk}^* (X_{gi} X_{hi} X_{rj} X_{sk} + X_{gi} X_{hj} X_{ri} X_{sk} + X_{gi} X_{hj} X_{rk} X_{si} \\
&\quad \quad + X_{gi} X_{hj} X_{rj} X_{sk} + X_{gi} X_{hj} X_{rk} X_{sj} + X_{gi} X_{hj} X_{rk} X_{sk}) \\
&\quad + \mu_\emptyset \sum^* X_{gi} X_{hj} X_{rk} X_{s\ell}.
\end{aligned}$$

Next, we write the terms of $n^4 \mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h \tilde{\beta}_r \tilde{\beta}_s)$ using \bar{X}_{ghrs} and similar moments.

The coefficient of μ_{4k} is $\sum_i X_{gi} X_{hi} X_{ri} X_{si} = n \bar{X}_{ghrs}$. The coefficient of μ_{3k} contains

$$\sum_{ij}^* X_{gi} X_{hi} X_{ri} X_{sj} = \sum_{ij} X_{gi} X_{hi} X_{ri} X_{sj} - \sum_i X_{gi} X_{hi} X_{ri} X_{si} = -n \bar{X}_{ghrs}$$

and after summing all four such terms, the coefficient is $-4n \bar{X}_{ghrs}$. The coefficient of μ_{2p} contains

$$\sum_{ij}^* X_{gi} X_{hi} X_{rj} X_{sj} = \sum_{ij} X_{gi} X_{hi} X_{rj} X_{sj} - \sum_i X_{gi} X_{hi} X_{ri} X_{si} = -n \bar{X}_{ghrs}$$

and accounting for all three terms yields $-3n \bar{X}_{ghrs}$.

The coefficient of μ_{1p} contains

$$\begin{aligned}
\sum_{ijk}^* X_{gi} X_{hi} X_{rj} X_{sk} &= \sum_{ijk} X_{gi} X_{hi} X_{rj} X_{sk} - \sum_{ij} X_{gi} X_{hi} X_{ri} X_{sj} \\
&\quad - \sum_{ik} X_{gi} X_{hi} X_{rj} X_{si} - \sum_{jk} X_{gi} X_{hi} X_{rj} X_{sj} + 2 \sum_i X_{gi} X_{hi} X_{ri} X_{si} \\
&= -n^2 \bar{X}_{gh} \bar{X}_{rs} + 2n \bar{X}_{ghrs}.
\end{aligned}$$

Summing all 6 terms, we find that the coefficient is

$$-2n^2 (\bar{X}_{gh} \bar{X}_{rs} + \bar{X}_{gr} \bar{X}_{hs} + \bar{X}_{gs} \bar{X}_{hr}) + 12n \bar{X}_{ghrs}.$$

The coefficient of μ_{\varnothing} is, using (7),

$$\begin{aligned} \sum_{ijkl}^* X_{gi} X_{hj} X_{rk} X_{s\ell} &= \sum_{ij} \left(X_{gi} X_{hj} X_{ri} X_{sj} + X_{gi} X_{hj} X_{rj} X_{si} + X_{gi} X_{hi} X_{rj} X_{sj} \right) \\ &\quad - 6 \sum_i X_{gi} X_{hi} X_{ri} X_{si} \\ &= n^2 (\bar{X}_{gh} \bar{X}_{rs} + \bar{X}_{gr} \bar{X}_{hs} + \bar{X}_{gs} \bar{X}_{hr}) - 6n \bar{X}_{ghrs}. \end{aligned}$$

We may summarize these results via

$$\mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h \tilde{\beta}_r \tilde{\beta}_s) = \begin{pmatrix} \mu_{4k} \\ \mu_{3k} \\ \mu_{2p} \\ \mu_{1p} \\ \mu_{\varnothing} \end{pmatrix}^T B \begin{pmatrix} \bar{X}_{ghrs}^*/n^2 \\ \bar{X}_{ghrs}/n^3 \end{pmatrix}, \quad \text{for } B = \begin{pmatrix} 0 & 1 \\ 0 & -4 \\ 1 & -3 \\ -2 & 12 \\ 1 & -6 \end{pmatrix},$$

where $\bar{X}_{gh,rs}^* = \bar{X}_{gh} \bar{X}_{rs} + \bar{X}_{gr} \bar{X}_{hs} + \bar{X}_{gs} \bar{X}_{hr}$, completing the proof of Lemma 2.

Appendix 3: moments of orthogonal random matrix elements.

We will need low order moments of orthogonal random matrices to study the moments of linear and quadratic test statistics under rotation sampling.

For integers $n \geq k \geq 1$, let $V_{n,k} = \{Q \in \mathbb{R}^{n \times k} \mid Q^T Q = I_k\}$, known as the Stiefel manifold. We will make use of the uniform distributions on $V_{n,k}$. There is a natural identification of $V_{n,1}$ with the unit sphere.

Let $Q \in \mathbb{R}^{n \times n}$ be a uniform random rotation matrix. This implies, among other things, that each column of Q is a uniform random point on the unit sphere in n dimensions.

By symmetry, we find that $\mathbb{E}(Q_{ij}) = 0$. Similarly $\mathbb{E}(Q_{ij}^2) = \mathbb{E}((1/n) \sum_{j=1}^n Q_{ij}^2) = 1/n$ and $\mathbb{E}(Q_{ij} Q_{rs}) = 0$ unless $i = r$ and $j = s$.

Anderson et al. (1987) give

$$\mathbb{E}(Q_{ij}^4) = \frac{3}{n(n+2)}. \quad (8)$$

We are interested in all fourth moments $\mathbb{E}(Q_{ij} Q_{k\ell} Q_{rs} Q_{tu})$ of Q . If any of j, ℓ, s, u appears exactly once then the fourth moment is 0 by symmetry. To see this, suppose that index ℓ appears exactly once. Now define the matrix \tilde{Q} with elements

$$\tilde{Q}_{ij} = \begin{cases} -Q_{ij} & j = \ell, \\ Q_{ij} & j \neq \ell. \end{cases}$$

If $Q \sim U(V_{n,n})$ then $\tilde{Q} \sim U(V_{n,n})$ too by invariance of $U(V_{n,n})$ to multiplication on the right by the orthogonal matrix $\text{diag}(1, 1, \dots, 1, -1, 1, \dots, 1)$, with a -1

in the j' 'th position. Then

$$\begin{aligned}\mathbb{E}(Q_{ij}Q_{kl}Q_{rs}Q_{tu}) &= \frac{1}{2}\mathbb{E}(Q_{ij}Q_{kl}Q_{rs}Q_{tu} + \tilde{Q}_{ij}\tilde{Q}_{kl}\tilde{Q}_{rs}\tilde{Q}_{tu}) \\ &= \frac{1}{2}\mathbb{E}(Q_{ij}Q_{kl}Q_{rs}Q_{tu} + Q_{ij}(-Q_{kl})Q_{rs}Q_{tu}) \\ &= 0.\end{aligned}$$

Similarly, because Q^\top is also uniformly distributed on $V_{n,n}$ we find that if any of i, k, r, t appear exactly once the moment is zero. If one index appears exactly three times, then some other moment must appear exactly once. As a result, the only nonzero fourth moments are products of squares and pure fourth moments. Their values are given in the Lemma below.

Lemma 4. *Let $Q \sim \mathcal{U}(V_{n,n})$. Then*

$$\mathbb{E}(Q_{ij}^2 Q_{rs}^2) = \begin{cases} \frac{3}{n(n+2)}, & i = r \text{ \& } j = s \\ \frac{1}{n(n+2)}, & 1_{i=r} + 1_{j=s} = 1 \\ \frac{n+1}{n(n-1)(n+2)}, & i \neq r \text{ \& } j \neq s. \end{cases}$$

Proof. The first case was given by Anderson et al. (1987).

For the second case, there is no loss of generality in computing $\mathbb{E}(Q_{11}^2 Q_{21}^2)$. The vector $(Q_{11}, Q_{21}, \dots, Q_{n1})$ is uniformly distributed on the sphere. Given Q_{11} , the point $(Q_{21}, Q_{31}, \dots, Q_{n1})$ is uniformly distributed on the $n-1$ dimensional sphere of radius $\sqrt{1 - Q_{11}^2}$. Therefore $\mathbb{E}(Q_{21}^2 | Q_{11}) = (1 - Q_{11}^2)/(n-1)$ and so

$$\mathbb{E}(Q_{11}^2 Q_{21}^2) = \frac{1}{n-1}\mathbb{E}(Q_{11}^2 - Q_{11}^4) = \frac{1}{n-1}\left(\frac{1}{n} - \frac{3}{n(n+2)}\right) = \frac{1}{n(n+2)}.$$

For the remaining case we let $\theta = \mathbb{E}(Q_{ij}^2 Q_{rs}^2)$ for $i \neq r$ and $j \neq s$. Summing over n^4 combinations of indices we find that

$$\sum_{i=1}^n \sum_{j=1}^n \sum_{r=1}^n \sum_{s=1}^n Q_{ij}^2 Q_{rs}^2 = \left(\sum_{ij} Q_{ij}^2\right)^2 = n^2$$

by orthogonality of Q . Therefore

$$n^2 = \mathbb{E}\left(\sum_{ij} \sum_{rs} Q_{ij}^2 Q_{rs}^2\right) = n^2 \mathbb{E}(Q_{11}^4) + 2n^2(n-1)\mathbb{E}(Q_{11}^2 Q_{12}^2) + n^2(n-1)^2 \theta.$$

Solving for θ we get

$$\theta = \frac{n^2 - \frac{3n}{n+2} - \frac{2n(n-1)}{n+2}}{n^2(n-1)^2} = \frac{n+1}{n(n-1)(n+2)}. \quad \square$$

Appendix 4: proof of Lemma 3.

Let $X_i \in \mathbb{R}^p$ where $p = |G|$ and $Y_i \in \mathbb{R}$ for $i = 1, \dots, n$. Both X_i and Y_i are centered: $\sum_i X_i = 0$ and $\sum_i Y_i = 0$.

The sample coefficients for genes $g \in G$ are given by the vector $\hat{\beta} = L = (1/n) \sum_i X_i Y_i$. The reference distribution is formed by sampling values of $\tilde{\beta} = (1/n) \sum_i X_i \tilde{Y}_i$ where \tilde{Y} is a rotated version of Y .

The rotation is one that preserves the mean of Y while rotating in the $n-1$ dimensional space of contrasts. As in Langsrud (2005), we let $W \in \mathbb{R}^{n \times (n-1)}$ be any fixed contrast matrix satisfying $W^\top W = I_{n-1}$ and $W^\top 1_n = 0_{n-1}$. Then the rotated version of Y is

$$\tilde{Y} = WQW^\top Y, \quad \text{where } Q \sim U(V_{n-1, n-1})$$

is a uniform random $n-1$ dimensional rotation matrix.

It is convenient to introduce centered quantities $X^c = W^\top X \in \mathbb{R}^{(n-1) \times p}$, $Y^c = W^\top Y \in \mathbb{R}^{n-1}$ and $\tilde{Y}^c = W^\top \tilde{Y} \in \mathbb{R}^{n-1}$. These sum to zero even when X , Y and \tilde{Y} do not. Their main difference from those variables is that they have $n-1$ rows, not n .

Now $\tilde{\beta} = (1/n) X^\top \tilde{Y} = (1/n) X^\top WQW^\top Y = (1/n) X^{c\top} Q Y^c$, so

$$\mathbb{E}(\tilde{\beta}) = (1/n) X^{c\top} \mathbb{E}(Q) Y^c = 0.$$

For the rest of the proof, we need the covariance matrix of $\tilde{\beta}$. Now

$$\mathbb{E}(\tilde{\beta} \tilde{\beta}^\top) = \frac{1}{n^2} X^{c\top} \mathbb{E}(Q^\top Y^c Y^{c\top} Q) X^c = \frac{1}{n^2} X^{c\top} \mathbb{E}(Q^\top Z Q) X^c$$

where $Z = Y^c Y^{c\top} \in \mathbb{R}^{(n-1) \times (n-1)}$.

The ij element of $Q^\top Z Q$ is $(Q^\top Z Q)_{ij} = \sum_{k=1}^{n-1} \sum_{\ell=1}^{n-1} Z_{k\ell} Q_{ki} Q_{\ell j}$ which has expected value

$$\sum_{k=1}^{n-1} \sum_{\ell=1}^{n-1} Z_{k\ell} 1_{k=\ell} 1_{i=j} / (n-1) = \frac{1_{i=j}}{n-1} \sum_{k=1}^{n-1} Z_{kk} = 1_{i=j} \frac{n}{n-1} \mu_2$$

where $\mu_2 = (1/n) \sum_{i=1}^n Y_i^2 = (1/n) \sum_{i=1}^n Y_i^{c2}$. That is

$$\mathbb{E}(Q^\top Z Q) = \frac{n\mu_2}{n-1} I_{n-1}$$

and so

$$\mathbb{E}(\tilde{\beta} \tilde{\beta}^\top) = \frac{\mu_2}{n(n-1)} X^{c\top} X^c.$$

In particular $\mathbb{E}(\tilde{\beta}_g \tilde{\beta}_h) = \mathbb{E}(\tilde{\beta} \tilde{\beta}^\top)_{gh} = \bar{X}_{gh} \mu_2 / (n-1)$, matching the value under permutation.

Appendix 5: cost analysis of $\text{var}(\tilde{C}_{G,w})$

Recall from Corollary 2 that in an experiment with $n \geq 4$ and genes g, h ,

$$\text{cov}(\tilde{\beta}_g^2, \tilde{\beta}_h^2) = \begin{pmatrix} \mu_2^2 \\ \mu_4 \end{pmatrix}^\top A^\top B \begin{pmatrix} \bar{X}_{gghh}^*/n^2 \\ \bar{X}_{gghh}/n^3 \end{pmatrix} - \frac{\mu_2^2}{(n-1)^2} \bar{X}_{gg} \bar{X}_{hh},$$

where $\bar{X}_{gghh}^* = \bar{X}_{gg} \bar{X}_{hh} + 2\bar{X}_{gh}^2$ and $A^\top B$ is a given 2×2 matrix.

To compute

$$\text{var}(\tilde{C}_{G,w}) = \sum_{g \in G} \sum_{h \in G} w_g w_h \text{cov}(\tilde{\beta}_g^2, \tilde{\beta}_h^2)$$

we need μ_2 , μ_4 and $A^\top B$ which are very inexpensive. We also need

$$S_1 \equiv \sum_{g \in G} \sum_{h \in G} w_g w_h \bar{X}_{gg} \bar{X}_{hh} = \left(\sum_{g \in G} w_g \bar{X}_{gg} \right)^2.$$

By expressing S_1 as a square, we find that it can be computed in $O(np)$ work, not $O(np^2)$ which a naive implementation would provide. We can compute all of the \bar{X}_{gg} 's in np multiplications and this is the largest part of the cost. If gene g belongs to many gene sets G we only need to compute \bar{X}_{gg} once and so the cost per additional gene set could be lower.

A similar analysis yields that

$$S_2 \equiv \sum_{g \in G} \sum_{h \in G} w_g w_h \bar{X}_{gghh} = \frac{1}{n} \sum_{i=1}^n \left(\sum_{g \in G} w_g X_{gi}^2 \right)^2$$

is also an $O(np)$ computation. Unfortunately $S_3 \equiv \sum_{g \in G} \sum_{h \in G} \bar{X}_{gh}^2$ does not reduce to an $O(np)$ computation. As written it costs $O(np^2)$. In cases where $p > n$, we can however reduce the cost to $O(n^2p)$ via

$$\begin{aligned} S_3 &= \sum_{g \in G} \sum_{h \in G} w_g w_h \left(\frac{1}{n} \sum_{i=1}^n X_{gi} X_{hi} \right)^2 = \frac{1}{n^2} \sum_{g \in G} \sum_{h \in G} w_g w_h \sum_{i=1}^n X_{gi} X_{hj} \sum_{j=1}^n X_{gj} X_{hj} \\ &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \left(\sum_{g \in G} w_g X_{gi} \right)^2. \end{aligned}$$

In terms of these sum quantities,

$$\text{var}(\tilde{C}_{G,w}) = \begin{pmatrix} \mu_2^2 \\ \mu_4 \end{pmatrix}^\top A^\top B \begin{pmatrix} (S_1 + 2S_3)/n^3 \\ S_2/n^3 \end{pmatrix} - \frac{\mu_2^2}{(n-1)^2} S_1.$$